



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/789,262

02/27/2004

Martin Deitch

CHA920030034US1

9375

23550

7590

09/19/2006

HOFFMAN WARNICK & D'ALESSANDRO, LLC
75 STATE STREET
14TH FLOOR
ALBANY, NY 12207

EXAMINER

PATEL, SHAMBHAVI K

ART UNIT

PAPER NUMBER

2128

DATE MAILED: 09/19/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/789,262

Applicant(s)

DEITCH, MARTIN

Examiner

Shambhavi Patel

Art Unit

2128

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☐ Responsive to communication(s) filed on 27 February 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-22 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-22 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 27 February 2004 is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
- ☐ Certified copies of the priority documents have been received.
 - ☐ Certified copies of the priority documents have been received in Application No. _____.
 - ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- 1) ☒ Notice of References Cited (PTO-892)
- 2) ☐ Notice of Draftsperson's Patent Drawing Review (PTO-948)
- 3) ☒ Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date 2/27/04.

- 4) ☐ Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____
- 5) ☐ Notice of Informal Patent Application
- 6) ☐ Other: _____

Art Unit: 2128

DETAILED ACTION

Claims 1-22 are pending.

Information Disclosure Statement

The information disclosure statement (IDS) submitted on 27 February 2004 is in compliance with the provisions of 37 CFR 1.97. Accordingly, the Examiner has considered the IDS as to the merits.

Claim Objections

In the instant application, claim 21 is dependent from claim 15. The Examiner believes that this is a mistake because there is no antecedent basis for this and asserts that Applicant intended for claim 21 to be dependent from claim 20.

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

1. Claims 1-22 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter. The Examiner asserts that the current state of the claim language is such that a reasonable interpretation of the claims would not result in any useful, concrete or tangible product. Regarding claims 1, 7, and 12, calculating a resource percentage time, a time slice percentage, a CP percentage, and then possibly causing the computer to not dispatch CPs does not produce a tangible output. Regarding claims 16, 17, 18, and 22, repeating the claimed steps until the observed consumption agrees with the defined consumption does not produce a tangible result.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

2. Claim(s) 1, 2, 7, 12, and 20-22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Rooney ('Intelligent Resource Director', 2002).

Regarding claim 1:

Rooney is directed to a method for modeling a behavior of an LPAR (logical partition) in a simulated computer operating in a time slice dispatch mode, comprising:

- a. beginning a next modeling interval (page 572 'State Sampling')
- b. calculating a resource percentage representing a percentage of total resources allocated to the LPAR (page 573 'Maximum Processor Demand' paragraph 2:
maximum_demand_percentage)
- c. calculating a time slice percentage for the LPAR based on the resource percentage (page 573 paragraph 1: *processor_using_samples*)
- d. determining a CP (central processor) percentage representing a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR (page 572 'State Sampling'). Four times a second, every work unit in the system is sampled, to learn *where each service class is spending its time and how much each class is using each resource.*

Rooney does not explicitly disclose comparing the CP percentage and the slice percentage to then accordingly allocate resources. However, a skilled artisan would knowingly have included this functionality because if the required resources (*slice percentage*) are greater than the available resources (*CP percentage*), then the CPs cannot be allocated to the LPAR.

Regarding claim 2:

Rooney is directed to the method of claim 1, including the further step of repeating each of the recited steps for a next modeling interval. (page 572 'State Sampling'). Four times a second, every work unit in the system is sampled, to learn *where each service class is spending its time and how much each class is using each resource*. Thus, the above process is repeated.

Regarding claim 7:

Rooney is directed to a tool for simulating operation of a computer having a system for modeling a behavior of an LPAR operating in a time slice dispatch mode, the modeling system comprising:

- a. a system for calculating a resource percentage representing a percentage of total resources allocated to the LPAR (page 573 'Maximum Processor Demand' paragraph 2: *maximum_demand_percentage*)
- b. a system for calculating a time slice percentage for the LPAR based on the resource percentage (page 573 paragraph 1: *processor_using_samples*)
- c. a system for determining a CP (central processor) percentage representing a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR (page 572 'State Sampling'). Four times a second, every work unit in the system is sampled, to learn *where each service class is spending its time and how much each class is using each resource*.

Rooney does not explicitly disclose comparing the CP percentage and the slice percentage to then accordingly allocate resources. However, a skilled artisan would knowingly have included this functionality because if the required resources (*slice percentage*) are greater than the available resources (*CP percentage*), then the CPs cannot be allocated to the LPAR.

Regarding claim 12:

Rooney is directed to a program product stored on a recordable medium for modeling a behavior of an LPAR in a simulated computer operating in a time slice dispatch mode, comprising:

- a. means for calculating a resource percentage representing a percentage of total resources allocated to the LPAR (page 573 'Maximum Processor Demand' paragraph 2: *maximum_demand_percentage*)
- b. means for calculating a time slice percentage for the LPAR based on the resource percentage (page 573 paragraph 1: *processor_using_samples*)
- c. means for determining a CP (central processor) percentage representing a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR (page 572 'State Sampling'). Four times a second, every work unit in the system is sampled, to learn *where each service class is spending its time and how much each class is using each resource*.

Rooney does not explicitly disclose comparing the CP percentage and the slice percentage to then accordingly allocate resources. However, a skilled artisan would knowingly have included this functionality because if the required resources (*slice percentage*) are greater than the available resources (*CP percentage*), then the CPs cannot be allocated to the LPAR.

Regarding claims 20-22:

Rooney is directed to a computer simulation tool for modeling LPAR behavior comprising:

- a. means for calculating a resource percentage representing a percentage of total resources allocated to the LPAR (page 573 'Maximum Processor Demand' paragraph 2: *maximum_demand_percentage*)
- b. means for calculating a time slice percentage for the LPAR based on the resource percentage (page 573 paragraph 1: *processor_using_samples*)
- c. means for determining a CP (central processor) percentage representing a percentage of time that all physical CPs in the computer being modeled have been allocated to the LPAR (page 572 'State Sampling'). Four times a second, every work unit in the system is sampled, to learn *where each service class is spending its time and how much each class is using each resource*.
- d. means for building a model for each LPAR specified in the computer simulation, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs (page 575 2nd paragraph)
- e. means for running each model and determining an observed consumption for each model (page 571 'WLM CPU weight-management configuration' 2nd paragraph *initial weight and current weight*)
- f. means for comparing the observed consumption with the defined consumption for all of the models (page 575 'Receiver Processing' 1st paragraph). The LPAR weight fix routine determines whether the receiver CPU delay bottleneck can be addressed by raising the partition processor weight (*consumption*).
- a. means for feeding back the observed consumption to the other models from each model that has an observed consumption that does not agree with the defined consumption (page 575 'Donor Selection'). After it has been determined that the weight

(consumption) of a service class needs to be increased, a donor whose weight (consumption) must be reduced as a result of this increase is selected.

- b. means for adjusting the defined consumption of each model based on the feedback (page 576 'Donor Projections' 1st paragraph). If a good trade is found, the partition weights of the receiver and donor are adjusted
- c. means for iteratively repeating the running, comparing, feeding and adjusting steps until the observed consumption agrees with the defined consumption for each model. (page 572 'Policy-adjustment framework'). This is repeated every ten seconds for every receiver class in need of resource allocation.

Rooney does not explicitly disclose comparing the CP percentage and the slice percentage to then accordingly allocate resources. However, a skilled artisan would knowingly have included this functionality because if the required resources (*slice percentage*) are greater than the available resources (*CP percentage*), then the CPs cannot be allocated to the LPAR.

3. Claims 3-6, 8-11, and 13-15 are rejected under 35 U.S.C. 103(a) as being unpatentable over Rooney ('Intelligent Resource Director', 2002) in view of Kyne ('z/OS Intelligent Resource Director', 2001).

Regarding claim 3:

Rooney does not explicitly disclose setting the resource percentage equal to 100% - a percentage of resources allocated to all other LPARs running in the simulated computer. Kyne teaches dividing the total amount of resources available (100%) among the LPARs running on the system (Kyne: page 61 table at bottom of page). At the time of the invention, it would have been obvious to one of ordinary skill in the art to combine the teachings of Rooney and Kyne because the method taught by Kyne provides

Art Unit: 2128

the ability to drive a processor at 100% while providing acceptable response times for the critical applications, and ensures that all resources are utilized by the right workloads (Kyne: page 4 1st paragraph).

Regarding claim 4:

The combination of Rooney and Kyne as applied to claim 3 above teaches basing the percentage of resources allocated to other LPARs on a weighting factor specified for each LPAR (Kyne: page 26 3rd paragraph), a number of logical CPs allocated to each LPAR (Kyne: page 48 4th paragraph), and a MIPS value for each LPAR (page 61 table at bottom of page).

Regarding claim 5:

The combination of Rooney and Kyne as applied to claim 4 above teaches the method of claim 4, wherein the MIPS value represents a maximum consumption that each LPAR could consume in an unrestrained processor (Kyne: page 65 2nd paragraph).

Regarding claim 6:

Rooney fails to explicitly disclose: $\text{time slice percentage} = (\text{resource percentage}) * (\# \text{ of physical CPs}) / (\# \text{ of logical CPs})$. Kyne teaches calculating the time slice percentage through the preceding equation (Kyne: page 55). At the time of the invention, it would have been obvious to one of ordinary skill in the art to combine the teachings of Rooney and Kyne because the method taught by Kyne provides the ability to drive a processor at 100% while providing acceptable response times for the critical applications, and ensures that all resources are utilized by the right workloads (Kyne: page 4 1st paragraph).

Art Unit: 2128

Regarding claim 8:

Rooney does not explicitly disclose setting the resource percentage equal to 100% - a percentage of resources allocated to all other LPARs running in the simulated computer. Kyne teaches dividing the total amount of resources available (100%) among the LPARs running on the system (**Kyne: page 61 table at bottom of page**). At the time of the invention, it would have been obvious to one of ordinary skill in the art to combine the teachings of Rooney and Kyne because the method taught by Kyne provides the ability to drive a processor at 100% while providing acceptable response times for the critical applications, and ensures that all resources are utilized by the right workloads (**Kyne: page 4 1st paragraph**).

Regarding claim 9:

The combination of Rooney and Kyne as applied to claim 8 above teaches basing the percentage of resources allocated to other LPARs on a weighting factor specified for each LPAR (**Kyne: page 26 3rd paragraph**), a number of logical CPs allocated to each LPAR (**Kyne: page 48 4th paragraph**), and a MIPS value for each LPAR (**page 61 table at bottom of page**).

Regarding claim 10:

The combination of Rooney and Kyne as applied to claim 9 above teaches the tool of claim 9, wherein the MIPS value represents a maximum consumption that each LPAR could consume in an unrestrained processor (**Kyne: page 65 2nd paragraph**).

Regarding claim 11:

Rooney does not explicitly disclose setting the resource percentage equal to 100% - a percentage of resources allocated to all other LPARs running in the simulated computer. Kyne teaches dividing the

Art Unit: 2128

total amount of resources available (100%) among the LPARs running on the system (**Kyne: page 61 table at bottom of page**). At the time of the invention, it would have been obvious to one of ordinary skill in the art to combine the teachings of Rooney and Kyne because the method taught by Kyne provides the ability to drive a processor at 100% while providing acceptable response times for the critical applications, and ensures that all resources are utilized by the right workloads (**Kyne: page 4 1st paragraph**).

Regarding claim 13:

Rooney does not explicitly disclose setting the resource percentage equal to 100% - a percentage of resources allocated to all other LPARs running in the simulated computer. Kyne teaches dividing the total amount of resources available (100%) among the LPARs running on the system (**Kyne: page 61 table at bottom of page**). At the time of the invention, it would have been obvious to one of ordinary skill in the art to combine the teachings of Rooney and Kyne because the method taught by Kyne provides the ability to drive a processor at 100% while providing acceptable response times for the critical applications, and ensures that all resources are utilized by the right workloads (**Kyne: page 4 1st paragraph**).

Regarding claim 14:

The combination of Rooney and Kyne as applied to claim 13 above teaches basing the percentage of resources allocated to other LPARs on a weighting factor specified for each LPAR (**Kyne: page 26 3rd paragraph**), a number of logical CPs allocated to each LPAR (**Kyne: page 48 4th paragraph**), and a MIPS value for each LPAR (**page 61 table at bottom of page**).

Art Unit: 2128

Regarding claim 15:

Rooney does not explicitly disclose setting the resource percentage equal to 100% - a percentage of resources allocated to all other LPARs running in the simulated computer. Kyne teaches dividing the total amount of resources available (100%) among the LPARs running on the system (**Kyne: page 61 table at bottom of page**). At the time of the invention, it would have been obvious to one of ordinary skill in the art to combine the teachings of Rooney and Kyne because the method taught by Kyne provides the ability to drive a processor at 100% while providing acceptable response times for the critical applications, and ensures that all resources are utilized by the right workloads (**Kyne: page 4 1st paragraph**).

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(b) the invention was patented or described in a printed publication in this or a foreign country or in public use or on sale in this country, more than one year prior to the date of application for patent in the United States.

4. Claim 16-19 are rejected under 35 U.S.C. 102(b) as being clearly anticipated by Rooney ('Intelligent Resource Director', 2002).

Regarding claim 16:

Rooney is directed to a method for modeling workload performance of a plurality of LPARs in a computer simulation, comprising

- a. providing a model for each LPAR specified in the computer simulation, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs (page 575 2nd paragraph)
- b. setting an initial defined consumption for each model, and running each model and determining an observed consumption for each model (page 571 'WLM CPU weight-management configuration' 2nd paragraph *initial weight* and *current weight*)
- c. comparing the observed consumption with the defined consumption for all of the models (page 575 'Receiver Processing' 1st paragraph). The LPAR weight fix routine determines whether the receiver CPU delay bottleneck can be addressed by raising the partition processor weight (*consumption*).
- d. for each model that has an observed consumption that does not agree with the defined consumption, feeding the observed consumption back to the other models (page 575 'Donor Selection'). After it has been determined that the weight (*consumption*) of a service class needs to be increased, a donor whose weight (*consumption*) must be reduced as a result of this increase is selected.
- e. adjusting the defined consumption of each model based on the feedback (page 576 'Donor Projections' 1st paragraph). If a good trade is found, the partition weights of the receiver and donor are adjusted.
- f. iteratively repeating the running, comparing, feeding and adjusting steps until the observed consumption agrees with the defined consumption for each model. (page 572 'Policy-adjustment framework'). This is repeated every ten seconds for every receiver class in need of resource allocation.

Regarding claim 17:

Rooney is directed to the method of claim 16, wherein the consumption is a measure of processor resources consumed by each LPAR (page 571 'WLM CPU weight-management configuration' 2nd paragraph).

Regarding claim 18:

Rooney is directed to a computer simulation tool for modeling workload performance of a plurality of LPARs in a computer simulation, comprising

- a. a system for building a model for each LPAR specified in the computer simulation, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs (page 575 2nd paragraph)
- b. a system for running each model and determining an observed consumption for each model (page 571 'WLM CPU weight-management configuration' 2nd paragraph *initial weight* and *current weight*)
- c. a system for comparing the observed consumption with the defined consumption for all of the models (page 575 'Receiver Processing' 1st paragraph). The LPAR weight fix routine determines whether the receiver CPU delay bottleneck can be addressed by raising the partition processor weight (*consumption*).
- d. a system for feeding back the observed consumption to the other models from each model that has an observed consumption that does not agree with the defined consumption (page 575 'Donor Selection'). After it has been determined that the weight (*consumption*) of a service class needs to be increased, a donor whose weight (*consumption*) must be reduced as a result of this increase is selected.

- e. a system for adjusting the defined consumption of each model based on the feedback (page 576 'Donor Projections' 1st paragraph). If a good trade is found, the partition weights of the receiver and donor are adjusted.
- f. a system for iteratively repeating the running, comparing, feeding and adjusting steps until the observed consumption agrees with the defined consumption for each model. (page 572 'Policy-adjustment framework'). This is repeated every ten seconds for every receiver class in need of resource allocation.

Regarding claim 19:

Rooney is directed to a program product stored on a recordable medium for modeling workload performance of a plurality of LPARs in a computer simulation, comprising

- d. means for building a model for each LPAR specified in the computer simulation, wherein each model includes a defined consumption that is dependent on a consumption of the other LPARs (page 575 2nd paragraph)
- e. means for running each model and determining an observed consumption for each model (page 571 'WLM CPU weight-management configuration' 2nd paragraph *initial weight* and *current weight*)
- f. means for comparing the observed consumption with the defined consumption for all of the models (page 575 'Receiver Processing' 1st paragraph). The LPAR weight fix routine determines whether the receiver CPU delay bottleneck can be addressed by raising the partition processor weight (*consumption*).
- g. means for feeding back the observed consumption to the other models from each model that has an observed consumption that does not agree with the defined consumption (page 575 'Donor Selection'). After it has been determined that the weight

(consumption) of a service class needs to be increased, a donor whose weight *(consumption)* must be reduced as a result of this increase is selected.

- h. means for adjusting the defined consumption of each model based on the feedback (page 576 **'Donor Projections'** 1st paragraph). If a good trade is found, the partition weights of the receiver and donor are adjusted
- i. means for iteratively repeating the running, comparing, feeding and adjusting steps until the observed consumption agrees with the defined consumption for each model. (page 572 **'Policy-adjustment framework'**). This is repeated every ten seconds for every receiver class in need of resource allocation.

Conclusion

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Shambhavi Patel whose telephone number is (571) 272-5877. The examiner can normally be reached on Monday-Friday, 8:00 am – 4:30 pm.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kamini Shah can be reached on (571) 272-2279. The fax phone number for the organization where this application or proceeding is assigned is (571) 273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).

SKP


KAMINI SHAH
SUPERVISORY PATENT EXAMINER